




**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

ESV



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.


In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated

25 OCT 1999



Patent Office

16 OCT 1998 03:07:55 - 003052
P01/7700 0.00 - 9822550.1

The Patent Office
Cardiff Road
Newport
Gwent NP9 1RH

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference

A25693

2. Patent application number
(The Patent Office will fill in this part)

15 OCT 1998

9822550.1

3. Full name, address and postcode of the or of each applicant (underline all surnames)

BRITISH TELECOMMUNICATIONS public limited company
81 NEWGATE STREET
LONDON, EC1A 7AJ, England
Registered in England: 1800000

Patents ADP number (if you know it)

1867002

If the applicant is a corporate body, give the country/state of its incorporation

UNITED KINGDOM

4. Title of the invention

COMPUTER COMMUNICATIONS

5. Name of your agent (if you have one)

ROBERTS, Simon Christopher

"Address for Service" in the United Kingdom to which all correspondence should be sent (including the postcode)

BT GROUP LEGAL SERVICES
INTELLECTUAL PROPERTY DEPARTMENT
HOLBORN CENTRE
120 HOLBORN
LONDON, EC1N 2TE

Patents ADP number (if you know it)

1867001

6. If you are declaring priority from one or more

Country

Priority application number
(if you know it)

Date of filing
(day/month/year)

earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day/month/year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

(See note (d))



9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description 10

Claim(s) 4

Abstract 1

Drawing(s) 2 (2 SHEETS, 2 FIGURES)

+ 2 (8)

10. If you are also filing any of the following, state how many against each item

Priority Documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)


Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature(s)

Date:



15 October 1998

ROBERTS, Simon Christopher, Authorised Signatory

12. Name and daytime telephone number of person to contact in the United Kingdom

Lisa Colton

0171 492 8146

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

This Page Blank (uspto)

COMPUTER COMMUNICATION

This invention relates to a method of communicating data from a server computer to a client computer.

5 Communication networks may be classified into two distinct groups: connection-based networks and connection-less networks. An example of a connection-based network is the public switched telephone network (PSTN) where a connection is established by the dialling of the required telephone number. An example of a connection-less network is the Internet, where data packets are
10 routed by the network to their destination, the user having no control over the route taken by the individual packets. The drawback of establishing a connection is that, generally, a high price must be paid in order to maintain the connection and thus guarantee delivery of the information, whereas the transmission of a stream of packets in a connection-less manner may fail due to changes in the intermediate
15 network elements, either due to equipment failure or network congestion. Disruption to the stream of packets does not cause such a problem if, for example, a text file is being downloaded from a server computer. However, if real-time data, such as video or audio, is to be transmitted then the disruption of a packet stream, so that packets arrive out of order or at a slower rate than is required for
20 data playback, may have serious repercussions. Accordingly, especially with regard to computer communications, there is a desire to attain the quality of service that is provided by connection-based transmission methods, without the additional cost that is incurred by setting up a connection.

 According to a first aspect of the present invention there is provided a
25 method of transmitting data from a server computer to a client computer over a communications network, the data being routed between the server and client computers by a network node; the network node having an input to receive data from the server computer, the input being connected to first and second buffer elements, said buffer elements being connected to an output channel of
30 predetermined bandwidth, wherein the first buffer element is preferentially allocated a portion of the output bandwidth and the second buffer element is allocated the remaining portion of the output bandwidth such that packets received in the first buffer element are transmitted in preference to packets received in the second buffer element; the method comprising the steps of:

(i) transmitting data from the server computer to the client computer using the first buffer element of the network node; and

(ii) upon receipt by the server computer of a first control signal from the client computer, transmitting data from the server computer to the client computer using the second buffer element of the network node. Preferably, the method comprises the further step of

(iii) reverting to transmitting data from the server computer to the client computer using the first buffer element of the network node upon receipt by the server computer of a second control signal from the client computer.

10 The first control signal may be generated by the client computer in response to the level of data stored in a data cache attaining a first, upper threshold value, whilst the second control signal may be generated by the client computer in response to the level of data stored in a data cache attaining a second threshold value which is lower than the first threshold value. The communications
15 route between the server computer and the client computer comprises more than one network node; and the selection of either the first or the second buffer elements in response to a control signal occurs within one or more of the network nodes which comprise the communications route between the server computer and the client computer.

20 A digital data carrier may contain computer executable code for loading into a computer for the performance of any of the above methods.

According to a second aspect of the present invention there is provided a method of receiving data at a client computer from a server computer, the data being routed over a communications network by a network node; the network
25 node having an input to receive data from the server computer, the input being connected to first and second buffer elements, said buffer elements being connected to an output channel of predetermined bandwidth, wherein the first buffer element is preferentially allocated a portion of the output bandwidth and the second buffer element is allocated the remaining portion of the output bandwidth
30 such that packets received in the first buffer element are transmitted in preference to packets received in the second buffer element; the method consisting the steps of:

(i) the client computer receiving data from the server computer via the first buffer element of the network node; and

(ii) the client computer receiving data from the server computer via the second buffer element of the network node in response to the transmission of a first control signal from the client computer to the server computer. The method may consist of the additional step of

- 5 (iii) the client computer receiving data from the server computer via the first buffer element of the network node in response to the transmission of a second control signal from the client computer to the server computer.

The first control signal may be generated by the client computer in response to the level of data stored in a data cache attaining a first, upper
10 threshold value, whilst the second control signal may be generated by the client computer in response to the level of data stored in a data cache attaining a second threshold value which is lower than the first threshold value. The communications route between the server computer and the client computer comprises more than
15 elements in response to a control signal occurs within one or more of the network nodes which comprise the communications route between the server computer and the client computer.

A digital data carrier may contain computer executable code for loading into a computer for the performance of any of the above methods.

- 20 According to a third aspect of the present invention there is provided a server computer for transmitting data to a client computer over a communications network, the data being routed by a network node, wherein the data is transmitted from the server as a plurality of data packets; the server computer in use transmitting data packets containing a first identifier to enable the preferential
25 forwarding of the data packets to the client computer at the network node; and wherein the server computer is responsive to a first control signal from the client computer to transmit data packets containing a second identifier to disable the preferential forwarding of the data packets to the client computer at the network node. Preferably, the server computer is additionally responsive to a second
30 control signal from the client computer to transmit data packets containing the first identifier to re-enable the preferential forwarding of the data packets to the client computer at the network node.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 shows a schematic depiction of a communications network over which a method of communication according to the invention may be used; and

Figure 2 shows a further schematic depiction of a communications network over which a method of communication according to the invention may be used.

Figure 1 shows a schematic depiction of a communications network over which a method of communication according to the invention may be used. Server computer 10 is connected to client computer 20 via a connection-less communications network 30, which includes at least one network node 40 in the communication route between the server computer 10 and the client computer 20. Each network node comprises two buffer elements 41 & 42, into which arriving packets are sorted on arrival at each network node. A flag in the header of each data packet determines which buffer element the packet is switched into. One of the buffer elements 41 is designated as a high priority buffer whilst the other buffer element 42 is designated as a low priority buffer. The high priority buffer 41 has preferential access to the output of the network node such that a majority of the bandwidth used by the network node is reserved for the high priority buffer. The remainder of the output bandwidth of the network node is reserved for the low priority buffer 42, although if the high priority buffer is not using its full bandwidth allocation then the low priority buffer may temporarily use the bandwidth allocated to the high priority buffer until there is a need for the high priority buffer to recommence use of its bandwidth allocation. An example of such a network node may be found in international patent application GB98/02727. The network node may be an IP (Internet Protocol) router, an ATM (Asynchronous Transfer Mode) switch which is switching IP data over an ATM network, or a device with similar functionality.

Communication is initiated by the client computer which sends a request to the server computer to transmit a data resource, such as a real time audio or video file, that is stored on the server computer. The request includes sufficient information about the client computer so that the server computer can transmit a stream of data packets to the client computer. A flag in the header of each transmitted data packet of the stream causes the stream of packets to be switched into the high priority buffer of each network node in the communications

path between the server computer and the client computer. This causes the data packets to be transmitted to the client computer with a guaranteed delivery rate.

The guaranteed delivery rate may be requested by the client computer when requesting the delivery of the real time data resource. Typically, the upper
5 limit to the guaranteed delivery rate will be determined according to the transmission rate available over the slowest communications link in the server computer-client computer communications path. This may be the link between the client computer and the communications network, which may be a modem dial-up link or a terminal adapter connecting a LAN to the communications network. The
10 lower limit to the guaranteed delivery rate may be determined according to the playback rate required for the real-time data resource, however, if the user of the client computer is prepared to wait before commencing playback of the data and has sufficient local cache capacity then the server may transmit data to the client at a rate which is lower than the playback rate, in response to a request from the
15 client. When delivering data using the high priority buffer element of each network node the guaranteed delivery rate should be significantly in excess of the minimum value of the guaranteed delivery rate. If the data is being delivered to the client via a proxy server, or a network cache held on a second server computer, then it is possible for the data stream to be delivered to the proxy server at a higher rate
20 than the client delivery rate, with the proxy server caching the additional data.

Under most circumstances, the stream of data packets is received by the client computer at a rate significantly greater than that required to play back the real-time data resource. The excess data received by the client computer must be stored in a local cache so that it is available for play back when required. This
25 local cache may be, for example, RAM or a disk drive. The size of the local cache is very much dependent upon the capabilities of the client computer, but there is a possibility that the amount of excess data received by the client computer is greater than the capacity of the local cache. In order to prevent overflow of the local cache, and subsequent loss of data, it is necessary for the client to send a
30 control message to the server computer when the local cache is filled to an upper threshold level, which should be slightly less than the total capacity of the local cache (allowing for the time need for the transmission and processing of the control message).

The control message, when received by the server computer, causes the server computer to instruct the internetworking protocols to change the flag that controls into which network node buffer element the data streams are switched, causing the data packets to be switched into the low priority buffer element in each network node in the communications path between the server and the client. This has the effect of stopping the guaranteed delivery of the data packet stream to the client, with the data packet stream being delivered on a 'best effort' basis. Additionally, the use of the low priority buffer means that packets may be lost between the server and the client, or that data packets may arrive out of sequence in the data packet stream. If these effects occur then the protocols that are used to control communications between the server and the client will have to be able to remedy them, i.e. by requesting re-transmission of packets and re-ordering packets received out of sequence.

If the client receives the packet stream at a rate which is still above the playback rate of the real-time data resource, then the local cache will continue to fill, albeit at a slower rate. This indicates a lightly loaded communications network, which tends to be a rare occurrence. If this situation were to continue, then there is a chance that the local cache could overflow, resulting in lost data. This situation may be remedied by the client sending a control signal to the server to request the server to transmit the data stream at a lower rate. Additionally, if there are low demands being placed on both the high priority buffer and the low priority buffer then there is a possibility that the data delivery rate could actually increase when the low priority buffer is being used. If this situation does occur then the client can send a control signal to the server to request the server to transmit the data stream at a lower rate.

If the packet stream is delivered to the client at exactly the playback rate of the real-time data resource then the local cache storage level will stay constant until all of the data packets have been received by the client, at which point the storage level of the local cache will decrease as the remainder of the real-time data resource is played back from the local cache.

If the packet stream is delivered to the client at less than the playback rate, then the local cache storage level will decrease at a rate equal to the difference between the playback rate and the data delivery rate. If the local cache is relatively large and the rate at which the local cache is emptied is relatively small

then the entire real-time data resource may be received and then played back by the client without the local cache emptying and data being lost. However, if the local cache is relatively small and the rate at which the local cache is emptied is relatively large, then the local cache will become empty whilst the data packet stream is being delivered at a rate lower than that required for playback, resulting in partial or complete loss of the playback of the real-time data resource. To prevent this, the client sends a second control message to the server when the storage level of the local cache reduces to a lower threshold level, which should be at least slightly above the point at which the local cache becomes empty. This second control message, when received by the server computer, causes the server computer to change the flag that controls into which network node buffer element the data packets are switched, causing the data packets to be switched into the high priority buffer element in each network node in the communications path between the server and the client. This has the effect of recommencing the guaranteed delivery of the data packet stream to the client, increasing the rate at which the data packet stream is received by the client and causing the client to receive sufficient data for uninterrupted playback of the real-time data resource whilst also increasing the storage level of the local cache.

This process may be repeated iteratively as required, decreasing the data delivery rate as the storage level of the local cache reaches the upper threshold and increasing the data delivery rate as the storage level of the local cache reaches the lower threshold, until the entire data resource has been delivered to the client and played back by the client.

The communications network used in the invention may use the standard internetworking protocol suite, i.e. TCP/IP, but the invention may be implemented using any connection-less communications network. The server computer may be a World Wide Web (WWW) server, whilst the client computer may be a personal computer (PC) or network computer (NC) running a WWW browser application program, such as Netscape Navigator™ or Microsoft Internet Explorer™.

It will be understood that in a large communications network there will be multiple server computers and multiple client computers, with many data streams being transmitted between various combinations of servers and clients. In order to ensure that the network nodes in the communications path between a particular server and a particular client are able to transmit the data stream at the guaranteed

data delivery rate it is necessary for some form of admission control to be applied to the network nodes. At the simplest level, this may involve each network node only admitting a certain number of data streams into the high priority buffer, with no further data streams being admitted once all of the bandwidth allocated to the high priority buffer is in use. The network node may also reject data streams if there is insufficient buffer capacity to support all of the data streams. Another consideration is that a data stream that is being transmitted by the low priority buffer may need to be switched to the high priority buffer if the local cache of the associated client computer is nearly empty. In order to ensure efficient network operation it is necessary that the probability that a data stream is not admitted into the high priority buffer from the low priority buffer is small.

An alternative variant to the above discussed method which can reduce the frequency with which data streams are switched from one buffer to the other can be achieved by the use of adaptive playback rates, the limits of which may be controlled by the client computer. If data is being transmitted to the client using the high priority buffers and the storage level of the local cache is increasing then it is likely that the upper threshold level will be reached, causing the client computer to signal to the server computer that the data should be transmitted using the low priority buffers. To prevent this occurrence and reduce the frequency with which data streams are switched between buffers, the client computer increases the playback rate based on an algorithm that is controlled by the instantaneous average rate at which the playback buffer is being filled. Provided the playback rate is not changed too rapidly and is kept within certain limits, the playback quality can be kept within acceptable limits. Similarly, if the local cache storage level is decreasing then the playback rate can be adaptively decreased in order to reduce the possibility of the lower threshold being reached. The playback rate limits can be determined in response to the transmission performance or the congestion of the network or by the user of the client computer. This technique may be particularly useful when multicasting because different users in the multicast group can experience different congestion conditions, and controlling the server under such conditions may be undesirable unless absolutely necessary.

This approach is primarily a method for temporarily delaying any action, other than within the client, for responding to transient congestion effects. The

advantages are that short duration congestion effects can be accommodated without having to change the buffer(s) being used to transmit data streams and that the number of control signals that might be generated in this period is reduced.

5 A further alternative method of transmitting real time data to the client can be envisaged where the request by the client to the server to initiate the transmission of the real time data resource contains additional instructions which cause defined portions of the real time data resource to be transmitted using selected buffers. For example, the request could instruct the server to transmit a
10 defined number of bytes using the high priority buffer and then start using the low priority buffer to transmit the remainder of the real time data resource. The client would use its knowledge of its own resources and operating speeds, current network performance, etc., to determine the number of bytes to be transmitted using the high priority buffer. This removes the need for the client to send a
15 control message to the server when the local cache storage level reaches the upper threshold. This approach could also be used when a return to the high priority buffer is triggered by the local cache storage level reducing to the predetermined lower threshold, the client requesting a portion of the real time data resource to be sent using the high priority buffer with the size of the portion being
20 calculated so as to restore the local cache storage level to approximately the upper threshold. An advantage of this method is that when returning to the use of the high priority buffer, the server would be able to specify to the network node (s) the number of bytes it would be transmitting using the high priority buffer, which may enable more versatile techniques for efficiently controlling use of the high priority
25 buffer.

It will be understood that the above methods of transmitting and receiving data can be implemented by the execution of a suitable computer program(s) on the server computer and the client computer. For example, if the server computer is a WWW server then such programs may be run in parallel with the software
30 program required to provide a WWW server, or integrated within such a WWW server program. Similarly, if the client computer is a PC or NC running a WWW browser then suitable program(s) may be integrated within or run in parallel with the WWW browser program. These suitable programs may be supplied on a data carrier such as a floppy disk, CD-ROM, magneto-optical disk, DVD (Digital Versatile

Disk), etc. Similarly, such programs could be provided over a telecommunications network. Figure 2 shows data carriers 70 and 80 which store suitable programs for the server computer and client computer respectively.

CLAIMS

1. A method of transmitting data from a server computer to a client computer over a communications network, the data being routed between the
5 server and client computers by a network node;

the network node having an input to receive data from the server computer, the input being connected to first and second buffer elements, said buffer elements being connected to an output channel of predetermined bandwidth, wherein the first buffer element is preferentially allocated a portion of
10 the output bandwidth and the second buffer element is allocated the remaining portion of the output bandwidth such that packets received in the first buffer element are transmitted in preference to packets received in the second buffer element; the method comprising the steps of:

(i) transmitting data from the server computer to the client computer
15 using the first buffer element of the network node; and

(ii) upon receipt by the server computer of a first control signal from the client computer, transmitting data from the server computer to the client computer using the second buffer element of the network node.

20 2. A method of transmitting data from a server computer to a client computer according to claim 1, wherein the method comprises the further step of

(iii) reverting to transmitting data from the server computer to the client computer using the first buffer element of the network node upon receipt by the server computer of a second control signal from the client computer.

25

3. A method of transmitting data from a server computer to a client computer according to claim 1 or claim 2, wherein the first control signal is generated by the client computer in response to the level of data stored in a data cache attaining a first, upper threshold value.

30

4. A method of transmitting data from a server computer to a client computer according to claim 2 or claim 3, wherein the second control signal is generated by the client computer in response to the level of data stored in a data

cache attaining a second threshold value which is lower than the first threshold value.

5. A method of transmitting data from a server computer to a client
5 computer according to any of claims 1 to 4, wherein:

the communications route between the server computer and the client computer comprises more than one network node; and

- the selection of either the first or the second buffer elements in response to a control signal occurs within one or more of the network nodes which comprise
10 the communications route between the server computer and the client computer.

6. A digital data carrier containing computer executable code for loading into a computer for the performance of any of claims 1 to 5.

- 15 7. A method of receiving data at a client computer from a server computer, the data being routed over a communications network by a network node;

the network node having an input to receive data from the server computer, the input being connected to first and second buffer elements, said buffer elements being connected to an output channel of predetermined
20 bandwidth, wherein the first buffer element is preferentially allocated a portion of the output bandwidth and the second buffer element is allocated the remaining portion of the output bandwidth such that packets received in the first buffer element are transmitted in preference to packets received in the second buffer element; the method consisting the steps of:

- 25 (i) the client computer receiving data from the server computer via the first buffer element of the network node; and

(ii) the client computer receiving data from the server computer via the second buffer element of the network node in response to the transmission of a first control signal from the client computer to the server computer.

30

8. A method of receiving data at a client computer from a server computer according to claim 7, wherein the method consists of the additional step of

(iii) the client computer receiving data from the server computer via the first buffer element of the network node in response to the transmission of a second control signal from the client computer to the server computer.

- 5 9. A method of receiving data at a client computer from a server computer according to claim 7 or claim 8, wherein the first control signal is generated by the client computer in response to the level of data stored in a data cache attaining a first, upper threshold value.
- 10 10. A method of receiving data at a client computer from a server computer according to claim 8 or claim 9, wherein the second control signal is generated by the client computer in response to the level of data stored in a data cache attaining a second threshold value which is lower than the first threshold value.
- 15 11. A method of receiving data at a client computer from a server computer according to any of claims 7 to 10, wherein:
the communications route between the server computer and the client computer comprises more than one network node; and
the selection of either the first or the second buffer elements in response
20 to a control signal occurs within one or more of the network nodes which comprise the communications route between the server computer and the client computer.
12. A digital data carrier containing computer executable code for loading into a computer for the performance of any of claims 7 to 11.
- 25 13. A server computer for transmitting data to a client computer over a communications network, the data being routed by a network node, wherein the data is transmitted from the server as a plurality of data packets;
the server computer in use transmitting data packets containing a first
30 identifier to enable the preferential forwarding of the data packets to the client computer at the network node; and
wherein the server computer is responsive to a first control signal from the client computer to transmit data packets containing a second identifier to disable

the preferential forwarding of the data packets to the client computer at the network node.

14. A server computer for transmitting data to a client computer according to
5 claim 11, wherein the server computer is additionally responsive to a second control signal from the client computer to transmit data packets containing the first identifier to re-enable the preferential forwarding of the data packets to the client computer at the network node.

ABSTRACT

Computer Communication

A method is provided for more efficiently transmitting data from a server computer
5 to a client computer over a connection-less network, the nodes of which are
capable of providing two classes of transmission quality. Data is transmitted using
the high priority class until a local cache at the client computer is filled to a certain
upper threshold, at which point the data is transmitted using the low priority class.
If the local cache at the client computer subsequently reaches a lower threshold
10 then the transmitted reverts to the use of the high priority class.

Figure 1




Figure 1




Figure 2

This Page Blank (uspto)